

SAP: an Architecture for Selectively Approximate Wireless Communication

Benjamin Ransford
Virta Laboratories, Inc.

Luis Ceze
University of Washington

Abstract

Integrity checking is ubiquitous in data networks, but not all network traffic needs integrity protection. Many applications can tolerate slightly damaged data while still working acceptably, trading accuracy versus efficiency to save time and energy. Such applications should be able to receive damaged data if they so desire. In today's network stacks, lower-layer integrity checks discard damaged data regardless of the application's wishes, violating the End-to-End Principle. This paper argues for *optional* integrity checking and gently redesigns a commodity network architecture to support integrity-unprotected data. Our scheme, called Selective Approximate Protocol (SAP), allows applications to coordinate multiple network layers to accept potentially damaged data. Unlike previous schemes that targeted video or media streaming, SAP is generic. SAP's improved throughput and decreased retransmission rate is a good match for applications in the domain of *approximate computing*.

Implemented atop WiFi as a case study, SAP works with existing physical layers and requires no hardware changes. SAP's benefits increase as channel conditions degrade. In tests of an error-tolerant file-transfer application over WiFi, SAP sped up transmission by about 30% on average.

1 Introduction

Today's predominant network protocols guarantee that data will be received exactly as it was sent, intact and usually in order. Even UDP, which allows datagram loss and reordering, typically runs on integrity-protected MAC and physical layers such as Wi-Fi.

This guarantee is necessary for applications that wish to abstract away communication details. However, as the population of endpoints shifts from reliable wired networks to error-prone wireless networks such as Wi-Fi and cellular, applications pay dearly for these convenient semantics. Bit errors result in relentless data retransmissions at the last hop, reducing the efficiency of chains of upstream links that must wait for last-hop physical and MAC-layer conditions to improve. Constant error recovery also increases endpoints' utilization of power-hungry wireless radios.

However, many applications deal with fundamentally noisy inputs and can tolerate errors. Computer vision, game mechanics, sensor data analysis, and search are all in this category. For these applications, low-layer integrity checking is often superfluous; in fact, lower layers violate the End-to-End Principle [33] by making judgments on the quality of network traffic. What is an error-tolerant application to do?

The solution to this mismatch is to put error handling under applications' control. If applications can claim possibly damaged data for themselves, rather than delegating integrity enforcement to lower layers, the network stack can provide better performance—crucially, lower latency and greater throughput—for error-tolerant applications, and both sides of a network connection can save precious energy on wireless radio use. An ideal realization of such a scheme would allow applications to transmit *some* important data—such as metadata—precisely while allowing errors in the rest.

This paper explores the consequences of providing *approximate* networking semantics to general-purpose applications. To this end, we propose, design, and evaluate Selective Approximate Protocol (SAP), an architecture for general-purpose approximate communication. SAP provides configurable integrity checking at multiple network layers, allowing applications to selectively compose traditional integrity-protected communication and higher-throughput unprotected communication. SAP builds on UDP-Lite, a transport protocol designed to accelerate multimedia streaming by relaxing transport-layer integrity checks [23]. SAP adds support for lower-layer integrity-check relaxation and a higher-level networking API for error-tolerant applications. SAP requires minor changes to applications' code and a small kernel patch under Linux, but it requires no hardware or physical-layer changes and is otherwise backward compatible with existing networks.

Contributions. This paper (1) extends the concept of general-purpose *approximate computing* to *selectively approximate networking*, in which error-tolerant applications can extend the benefits of approximation beyond individual machines; and (2) defines and evaluates SAP, a cross-layer embodiment of selectively approximate networking for off-the-shelf Wi-Fi devices.

Approximate network semantics offer several advan-

tages for applications. First, an application that tolerates errors can increase its throughput by accepting data more quickly; application-specific *semantic* integrity checks can replace checksums. Second, under adverse network conditions, an application can avoid unpredictable lag times due to retransmissions that disrupt application quality. Third, an application can achieve greater usable range [31] or tolerate more interference for a given output quality. Finally, an application can save energy by reducing the number of bytes transmitted.

Unlike past approaches to error-tolerant communications that have focused on a single class of applications (e.g., streaming media [37, 19]), SAP provides a *general* mechanism that is format agnostic and usable by any application. The SAP prototype is implemented atop IP and 802.11, but its key properties are portable to other stacks.

2 Motivation and Background

As wireless networks replace wired networks, transmission errors become more burdensome. Errors at the MAC layer trigger retransmissions that decrease goodput and channel availability. Errors become more pronounced as transmission rates and distances increase. These problems are generally not present—or negligibly rare—in wired networks.

An unfortunate side effect of MAC-layer errors is that they have a congestion-like effect on networks upstream. To a sender several hops away, MAC-layer retransmissions at the last hop manifest as delays in reaching the receiver, which both sides incorrectly interpret as congestion, causing retransmissions along the entire path. Last-hop errors thus cause slowdowns along paths that are *mostly* reliable.

If slightly damaged payloads did not require retransmission, applications could experience lower latency and greater goodput, and the reduction in wireless delays could alleviate upstream congestion and reduce MAC-layer wireless contention. The key question is whether damaged payloads can still be useful to applications.

Characterizing wireless errors. A simple experiment in a noisy WiFi environment—described fully in Section 4—yields a surprising result: even at high rates of frame retransmission, the bit error rate (BER) in frames considered corrupt remains below a threshold that would be tolerable to many applications. Furthermore, BER does *not* scale directly with the fraction of frames that are retransmitted, suggesting that, on average, a small number of bad bits cause a large fraction of retransmissions.

This paper proposes Selective Approximate Protocol (SAP) as a response to retransmission behavior that is inappropriate for error-tolerant applications. SAP allows applications to choose approximate or precise network semantics at run time to suit their needs. Error-tolerant

applications like those studied in previous literature on approximate computing (e.g., image rendering, game-play updates) can choose *approximate* transmission for information that is noisy or contains uncertainty. Applications that already reduce resource use via approximate computing can match their communication strategies to their computation strategies. The main goal of SAP is to reduce the time and energy that applications must spend on communication.

2.1 Background: Data Integrity

The OSI layered networking model defines a stack of abstraction layers that enable progressively simpler communication primitives. The bottom layer encapsulates the physical properties of the network, e.g., radio transmission encoding, and the top layer encapsulates applications’ communication via straightforward *send* and *receive* operations. This model delegates responsibility for well-formedness to each layer. This section focuses on 802.11 (Wi-Fi) networks running IPv4, but many of the terms are applicable to other wired and wireless networks, *mutatis mutandis*.

In 802.11 networks, *stations* transmit *frames* via a decentralized Media Access Control (MAC) protocol that coordinates access to the wireless channel. Each frame includes a mandatory 32-bit *frame check sequence* (FCS) that is a cyclic redundancy check (CRC) of the rest of the frame. The standard does not mandate a specific action for stations to take when an FCS check fails; instead, it says that “All STAs shall *be able* to validate every received frame using the frame check sequence” [18, §8.1] (emphasis added) and it mandates that stations will issue *positive acknowledgments* (ACKs) upon correctly receiving certain types of frames [18, §9.3.2.2]. Absent an ACK, senders retransmit frames up to a retry limit.

At the network layer, IPv4 packets carried in 802.11 frames include another 16-bit checksum that covers only the IP header [30]. The IPv4 standard mandates that receivers “silently discard every datagram that has a bad checksum.” [30, §3.1]

At the transport layer, TCP and UDP include packet (datagram) checksums [6]. The 16-bit TCP checksum is required for TCP packets; receivers must silently discard packets that fail the CRC check. UDP, however, allows senders to fill a datagram header’s 16-bit checksum field with zero bits, which tells the receiver not to compute a checksum over the datagram, and which may save checksum-computation time for the sender and receiver.

The net effect of these integrity checks across layers is that *multiple independent mechanisms can cause damaged data to be discarded*. If the link layer did not include a checksum, then transport-layer checks would be the only line of defense against corruption. In mod-

ern networks that include FCS-like low-level integrity checks (e.g., Ethernet and 802.11), higher-level checks are largely redundant. In response, IPv6 omits header checksums because they are redundant with lower- and higher-layer checksums meant to protect against bit errors and format problems [8].

For the protection they provide, checksums are not perfect; they are inadequate on their own to provide perfect data integrity. Checksums *do* offer receivers some protection against malformed data from ill-behaved network stacks, but the simple CRCs on network packets offer imperfect, skewed output distributions and cannot detect certain kinds of splices and other small changes [40]. Consequently, the traditional advice to application designers is to perform application-level checks once data arrives. Application-level integrity checks thus represent *another* layer of protection—one that is arguably more important than the lower-level checks because it can incorporate arbitrary decision-making logic and quality control. With respect to the End-to-End Principle [33], the application-layer check is the most important.

Relaxing integrity checks. The UDP-Lite variant of UDP [23, 24] exists to provide *configurable* checksum protection of UDP packets for applications that can tolerate some errors. UDP-Lite replaces the UDP header’s *Length* field (which can be inferred from the enclosing IP datagram’s *Length* field) with a *Checksum Coverage* (*cscov*) field that specifies the number of octets in the UDP-Lite datagram to be included in a checksum computation. A value of 0 in the checksum coverage field indicates full coverage, equivalent to a UDP checksum. The UDP-Lite header is always included in the checksum computation. As in UDP, datagrams with errors in the *cscov*-covered range are unceremoniously dropped.

Applications can control UDP-Lite checksum coverage via *setsockopt* calls at run time. An application can provide a *cscov* value so that none, all, or the first *n* bytes of a payload are covered by the checksum.

Past work has shown *in simulation* that UDP-Lite can decrease loss rates for multimedia applications by deeming partially damaged data acceptable [15, 27]. Like link-layer coding techniques that implement packet and header correction from redundant information [26], these techniques’ tolerance of damaged payloads can mostly preserve application-layer quality metrics while reducing retransmission rates.

SAP uses UDP-Lite to provide configurable *approximate* semantics to general applications beyond multimedia, offering approximate communication as a complementary mode to *precise* guaranteed in-order delivery. SAP combines these transport-layer mechanisms with the MAC-layer mechanisms necessary to preserve the semantics across the network stack, embracing the end-to-end principle [33] to allow applications on both ends of



Figure 1: A JPEG image transmitted in SAP’s precise (left) and approximate (right) modes showing that bit errors render the image degraded but recognizable. (Photo of a red-necked wallaby [*Macropus rufogriseus*], © 2002 California Academy of Sciences.)

a communication link to fully control data delivery.

2.2 Background: Approximate Computing

The growing body of work in approximate computing observes that many applications can save time and energy by engaging fewer resources—whether by skipping work [17], reducing voltage [10], storing information in less reliable memories [25, 35], or training neural networks to emulate expensive computation [11].

Applications that benefit from approximate computing include 3D rendering and game engines [34], database query processing [2, 1], computer vision and robotics [11], sensor data storage and analysis [35], and weather simulations [9].

To date, approximate computing’s performance and efficiency improvements have largely been confined to *local* speedups on a single device. Unfortunately, energy-constrained devices that benefit from these speedups, such as phones or embedded devices, can waste all of these performance gains on data transmission—even when the data’s recipients do not require perfect fidelity. Reliable TCP/IP stacks on reliable MAC layers are imperfectly matched with these application domains.

Previous work on approximate computing has not explicitly sought to optimize *communication* costs, leaving an opportunity for further savings. Radios are major consumers of power in modern computers. On a mobile phone, the WiFi and cellular radios require an order of magnitude more power than even the CPU or memory [7]. Further, transmitting data requires not only the radio but the CPU to be awake, since network stacks are predominantly implemented in software.

Even heavily encoded formats, such as JPEG, are often robust against a certain amount of error (Figure 1) [13]. The intuition behind this sur-

prising robustness is that the information that influences human perception—such as low-frequency image components—often comprise only a small part of the total image data, with potentially inconsequential details (e.g., high-frequency components or color tables) comprising most of the bits. On one hand, a randomly placed bit error is more likely to affect inconsequential data; on the other hand, unfortunately located bit errors can completely hamper decodability. The necessary division of data between precise and approximate handling, proposed for EnerJ [34], can also apply to data formats being transmitted: precise data can receive checksum coverage, and approximate data can be transmitted without checksum coverage.

Of course, certain kinds of data are fundamentally precise and are incompatible with approximate transmission. Encrypted or heavily compressed payloads are prime examples of such precise data. Still, a great deal of Internet and LAN traffic occurs without encryption or compression [36], so mechanisms to relax integrity requirements still apply.

3 Design of SAP

SAP is a transport protocol and API designed to expose potentially damaged data to applications that want it. SAP’s cross-layer design philosophy can be summarized as follows:

- Integrity checking should be under *applications’* control; applications should be able to receive potentially damaged data if they wish.
- Applications should be allowed to switch between precise and approximate communication modes quickly and without special privileges.
- For reasons of cost and inertia, approximate communication should be possible on unmodified hardware and physical layers and should travel concurrently with conventional traffic.

This design philosophy underlies the *approximate networking model*, which we define as follows:

1. A unit of data is either *precise* or *approximate*.
2. Precise data will have accompanying metadata (e.g., a checksum) that the receiver can use to confirm exact reception.
3. Approximate data can optionally include similar metadata, but sub-application-layer mechanisms at the receiver will not use this metadata to decide whether to pass the data upward.
4. A single transmission can include both precise and approximate data, provided it includes metadata that unambiguously marks the two kinds.
5. Precise data requires acknowledgments from receivers and is guaranteed to arrive in order; approximate data does not.

6. Control data (e.g., acknowledgments, connection setup, and teardown) are precise.

SAP embodies the approximate networking model by building it atop 802.11 [18] and UDP-Lite [23]. It necessarily involves multiple layers for the reasons outlined in Section 2.1, and maps to the model as follows:

1. SAP provides sockets that are precise by default, and an API to switch the socket to approximate mode and back.
2. Sockets are built on UDP-Lite datagrams that offer partial or full checksum coverage of precise data.
3. SAP adds an optional *approximate* mode to 802.11 to carry approximate data, selectively changing retry behavior.
4. A checksum coverage field inherited from UDP-Lite demarcates where precise data ends and approximate data begins within a datagram.
5. Receivers send (precise) acknowledgments for precise data; senders re-send precise data until they receive acknowledgments.
6. SAP always uses precise transmissions for control data including connection setup and teardown; it implements a simple TCP-like handshake.

The remainder of this section details the design of SAP with reference to conventional wireless network communication using 802.11 and TCP/IP. Section 3.4 gives lower-level implementation details.

3.1 Changes to 802.11

SAP makes only one small software change to 802.11 transmission: the 802.11 driver examines outgoing payloads—specifically, inspecting `sk_buff` structures for UDP-Lite headers that indicate partial checksum coverage—to detect whether the application has marked the data as approximate. If so, the driver sets the 802.11 retry limit to zero to disable retransmission of the frame if it does not receive an acknowledgment frame.

Receivers are subject to a larger number of software changes to support approximate data. SAP provides a boolean driver-level switch that controls both hardware checksum offloading (disabled when the switch is in approximate mode) and the kernel’s decision procedure that leads to each frame’s acceptance or rejection. When the switch is in the *approximate* position, the 802.11 driver processes *all* frames instead of simply dropping those that fail the FCS check. The MAC-layer logic is left intact, so the driver will, for example, still ignore frames destined for other stations (which may include frames with corrupted address fields).

With the driver-level mode switch in the *precise* position, 802.11 behavior is unchanged from conventional 802.11 transmission and reception.

3.2 Changes to Applications

Because it offers its own socket API, SAP requires changes to application code. The API presents *send* and *receive* functions that map exactly onto the familiar sending and receiving functions of conventional sockets. The behavior of these functions depends on the socket's *approximation mode*, which is *precise* at socket creation but can switch between *precise* and *approximate* at run time. The socket's approximation mode is independent of the aforementioned kernel-level switch.

Applications that use UDP, such as media-streaming applications, are trivially portable to SAP, especially if they do not rely on guaranteed in-order delivery. Since SAP sockets wrap UDP sockets, an application that already uses datagram-oriented communication need only substitute calls to the SAP versions of the socket functions, and optionally change the approximation mode based on application requirements.

Porting TCP applications is only slightly more challenging. As a proof-of-concept embodiment of the approximate networking model, SAP reimplements only the most important property of TCP on which applications rely: guaranteed in-order delivery for precise data. It does so via the same mechanism that TCP uses, namely ACKs, though it sends them as separate (precise) transmissions rather than header fields. The more significant difference is that SAP, based on UDP, is datagram oriented rather than stream oriented; application code must be more explicit about when to send and receive data. The same problem affects any port of a TCP-based application to UDP, so we do not belabor it here. Section 4 describes porting a TCP-based application to SAP.

Regardless of any changes to socket code, applications using SAP become responsible for data integrity. In contrast to cheap but imperfect checksums, applications can implement semantic, application-aware integrity checks (e.g., “is this data point within n units of the previous data point?”) when they accept approximate, potentially damaged data.

3.3 Design Implications

According to the standard, an 802.11 transmitter is allowed to adjust its bitrate, maximum transfer unit (MTU), selected channel, transmit power, retry limit, request-to-send (RTS) threshold, and power-management options. Retransmission rate, as a proxy for channel quality, is one metric that governs how aggressively a transmitter adjusts these parameters. Each of these causes performance to degrade under poor channel conditions, so stations often overprovision for current conditions. SAP's introduction of error tolerance enables stations to adjust these parameters more aggressively, re-

ducing overprovisioning.

Under SAP, an 802.11 station can: reduce its transmit power (which makes decoding errors more likely by reducing the signal-to-noise ratio); increase its bitrate (which makes decoding errors more likely by making symbols “smaller”); increase its RTS threshold (which makes bit errors in frames more likely by gambling that longer transmissions will be error free); and decrease its retry limit as described above (since each try is more likely to succeed). At the MAC layer, the operating system can increase the maximum transfer unit (MTU), effectively gambling on longer transmissions in the same way as increasing the 802.11 driver's RTS threshold.

3.4 Implementation

SAP comprises two parts: a small set of device driver and software MAC layer modifications and a user-space library that implements SAP's user-space protocol components. We implemented the kernel modifications against version 3.8.13 of the Linux kernel.

Our changes to the Linux kernel are: 40 additional lines of code in the `mac80211` subsystem, which implements the 802.11 MAC layer, to modify its frame-dropping logic; and 165 new lines of code in the `ath5k` 802.11 card driver, to modify its frame handling and implement SAP's approximation mode switch.

The SAP user-space library, `libsap`, implements an API atop UDP-Lite sockets and comprises 766 lines of C code. A `sap_sock_t` socket object keeps metadata to track acknowledgment numbers for precise datagrams and the socket's connection state; there is no global state. An application can act as sender, receiver, or both with the appropriate API calls.

For senders, three API functions handle connection setup and teardown. The `sap_connect` function stores endpoint information in a socket object so it is retained between datagrams. `sap_connect` then sends a precise SAP_PING message to the receiver and waits for an acknowledgment. (If no acknowledgment arrives within a timeout period, the `sap_connect` call returns an error; it is ideal to set up the receiver to listen before the sender issues a ping.) After the receiver acknowledges the ping, the socket is in the connected state and the sender is free to send datagrams. When it is done with a SAP transmission (e.g., an approximate file transfer), the sender closes its end of the socket by issuing a precise SAP_FIN message to the receiver. The sender waits up to a timeout period for the receiver to acknowledge the SAP_FIN message, then calls `sap_close` which simply frees the socket's memory. (It is wise for senders to use multiple threads of execution so that precise transmissions, such as SAP_FIN, do not block traffic for other connections.)

Receivers use `sap_listen` to set up a SAP socket,

which implements the other side of the above protocol on top of UDP-Lite. For brevity, we omit the details of the receive side. The salient property of receivers is that they handle precise and approximate traffic differently.

4 Evaluation

This section evaluates the degree to which SAP meets its design goals—in short, how approximate transmission of data affects applications’ performance–quality balance.

The evaluation studies three applications, each time focusing on a different aspect of SAP’s behavior: a simple synthetic application that streams precomputed values to a receiver over an approximate socket; a web server (lighttpd) that streams files to clients in precise or approximate mode; and a location-tracking application that emits periodic GPS readings for speed calculation.

We evaluated SAP between PCs in a multistory office environment with interference from many nearby enterprise-grade access points. For SAP senders and receivers, we used Dell Optiplex GX620 workstations running Ubuntu Linux 12.04.5 LTS and kernel version 3.8.13 (patched as described in Section 3.4). Each workstation used a TP-Link WN350GD 802.11b/g PCI Express card controlled by the ath5k kernel driver. When testing transmission between two nodes using SAP, we kept one node stationary in a $3\text{ m} \times 4\text{ m}$ office and moved the other node to different points in the building, as close as 1 m and as far away as 12.2 m.

4.1 Applications

Streamer. We designed the first application, *Streamer*, to characterize 802.11 frame errors and their effect on application-level data delivery. Streamer sends a fixed number of datagrams to a receiver exclusively in approximate mode. The receiver captures and analyzes all data that reaches the receiving application, and it also collects and analyzes 802.11 statistics from the Linux kernel.

The sender and receiver agree on the contents of the entire stream in advance by choosing an initial 32-bit seed, which the sender sends to the receiver precisely, and which is incremented for each datagram. Each datagram consists of the result of passing the 32-bit counter through a hash function that outputs 32 unbiased bits.

Lighttpd. We chose a web server application to evaluate (1) the difficulty of porting a TCP-based application to SAP and (2) the throughput performance of SAP’s precise and approximate modes versus precise TCP transmission. *Lighttpd* [22] is a popular event-based web server. Web servers are especially good candidates for SAP’s *selectively* approximate communication because they serve a variety of file types via the ubiquitous HTTP protocol. A typical webpage visit from a browser loads

several resources—HTML pages, scripts, and images—often from the same server, but only some of these data types are damage tolerant. We modified lighttpd to serve content via SAP as follows.

HTTP requests proceed as normal, with clients initiating connections to the web server via TCP and issuing requests. A SAP-aware client adds two headers: X-SAP-Approx followed by a list of MIME types for which the client can tolerate errors (e.g., image/jpeg); and X-SAP-Port, which tells the server on which UDP port the client will listen for a SAP connection.

Upon receiving an HTTP request with SAP headers, the modified lighttpd loads files from backing storage and determines their MIME types as usual, but it also matches MIME types against those that the client requested to receive approximately. In the case of a match, the server initiates a SAP connection back to the client on the client’s requested port, sets the SAP socket’s mode to *approximate*, and sends the contents of the file in 1 KB datagrams. Files whose MIME types do not match the client’s X-SAP-Approx header are sent via TCP.

For a direct comparison of precise and approximate transmission, we also added an X-SAP-Force-Precise HTTP request header to tell lighttpd to send its responses via SAP in precise mode instead of approximate mode.

A set of concurrent HTTP requests may cause the server to initiate multiple SAP connections to the client, but for simplicity, SAP includes no multiplexing facility. To regain the parallelism that multiple TCP streams offers, clients can issue each HTTP request with a different X-SAP-Port value, then `sap_listen` on all of the corresponding ports at once.

These changes added 178 lines of C code to lighttpd version 1.4’s 54632 lines and required only a few hours of work for a single programmer, most of which time went to understanding lighttpd’s request workflow.

Tracker. To evaluate SAP’s effect on the *quality* of transmitted data, we built a location-tracking application with a straightforward quality metric. In practice, quality metrics are application specific. Applications that use approximate communication with SAP should implement their own quality metrics that perform sanity checking and damage control (e.g., requesting retransmission of unacceptably damaged pieces of data).

Tracker is a simple location tracker that emits a latitude and longitude (a pair of 64-bit double-precision floating point values) at a predefined rate. A sender emits these values one at a time and a receiver calculates the cumulative moving average of the sensor’s speed with each update, finally emitting a final moving average covering the duration of the trace. We used a trace of GPS readings collected by a mobile phone while its owner walked around a campus. To smooth the inevitably noisy GPS signal, we followed the example of previ-




Figure 2: The fraction of frames that must be retransmitted increases dramatically with increasing 802.11b/g bitrate and distance.

ous approximate-computing work [5] and made the receiver discard location points that would have indicated foot speed outside a reasonable bound—in this case, any speed greater than the maximum running speed of an Olympic sprinting medalist.

In a real deployment of this application, a moving node associated with a WiFi network (via one or many access points) would experience highly variable channel quality [4], making continuous reliable transmission over TCP difficult. The SAP version of Tracker aims to constrain tracking-data error metrics while providing steady transmission that is not fraught with delays. We also implemented a TCP version that operates analogously.

The SAP version of Tracker comprises 152 lines of C for the listener, 101 lines of C for the sender, and 34 lines of shared C code for dealing with GPS measurements. The TCP version comprises an extra 50 lines of code, mostly for buffer management.

4.2 Characterizing Frame Retransmission Behavior

We used the *Streamer* application to transmit predictable, unbiased bit patterns in order to understand two factors influencing SAP’s performance: the proportion and nature of retransmitted frames, and how well SAP might be able to recover usable information for receivers that can tolerate damage. All tests were performed during the daytime in a busy office building with many enterprise access points nearby. When we varied the distance between nodes, distances over 5m meant the moving node was in nearby rooms on the same floor (i.e., transmissions traversed walls). We set up an ad hoc 802.11 network with an invisible SSID for SAP nodes.

To measure rates of frame retransmission in an 802.11 network, we sent 10MB of predictable data between two of our test machines using the *Streamer* application, varying the location of one of the nodes and testing five times per location (50MB total). We used SAP in *pre-*




Figure 3: Bit error rate does not scale directly with the fraction of corrupt frames. At higher bitrates, despite a greater fraction of frames being corrupt, coding gains result in decreasing BER in damaged frames.

cise mode so that the sender would retransmit frames. Figure 2 plots the fraction of frames that the sender retransmitted at least once because the receiver, seeing they did not pass the FCS check, did not acknowledge them. The salient feature of Figure 2 is that distance and bitrate have a deleterious effect on correct frame transmission, with retransmission rates exceeding 30% on average when the nodes used the highest bitrate at 12.2m.

A second question is: if frames were to be accepted despite FCS failures, how bad would the damage be, from an application’s perspective? To answer this question, we sent 10MB in SAP’s *approximate* mode via the *Streamer* application and measured the bit error rate (BER) of the payloads that arrived in damaged frames. We varied the distance between the sender and receiver and performed five trials at each distance as in the previous experiment. Figure 3 plots BER versus bitrate and distance for this experiment. Even at distances and bitrates with a high proportion of damaged frames, the bit error rate remained low. Most notably, at 48 and 54Mbps, coding gains drive the BER well below 5% in damaged frames. Counting the non-erroneous bits in damaged frames along with all the bits in correctly received frames, transmissions at the highest bitrates consisted of 99.6% correct bits—an error rate that is tolerable to many applications [34, 2].

Finally, to forecast the effect of accepting damaged data, we measured the frame loss rate (FLR) in the *Streamer* application several ways. At all the positions described in the previous experiment, we transmitted 10MB in SAP’s *approximate* mode and measured: (1) the fraction of frames that did not arrive at the receiver, and (2) the fraction of frames that arrived with damage (i.e., that failed their FCS checks). Without SAP’s modifications to the Linux kernel, the effective frame loss rate would have been the sum of the missing and damaged frame rates. With SAP’s modifications, the effective frame loss rate reflects only the fraction of *lost* frames




Figure 4: Accepting damaged frames (FLR/B, dotted blue line) with SAP allows a receiver to decrease the effective frame loss rate (FLR, solid red line).

(which we call FLR/B), since SAP and the Streamer application make no attempt to recover these. Figure 4 shows the rates of payload corruption and loss over all bitrates and distances. Note that packet loss may *also* be acceptable to applications that can accept damaged data, suggesting that a nonzero frame loss rate is reasonable.

4.3 Web Throughput

To evaluate SAP’s effect on throughput for a complete application (lighttpd), we repeatedly sent large compressed (JPEG) images from a sender to a receiver, both nodes on an ad-hoc network of the machines described earlier. Nodes were connected to the building network via wired Ethernet, which we used as a backchannel to coordinate the sender and receiver’s configurations between experiments. Lighttpd listened on both the Ethernet and 802.11 interfaces.

Throughput versus channel quality. We used the distance between sender and receiver as a proxy for channel quality, and measured how long it took to transfer an identical file in several different ways: via (precise) TCP, precise SAP, and approximate SAP.

In this experiment, the receiver first warmed the sender’s cache by requesting the file over TCP via the wired Ethernet interface. The receiver then issued sequential HTTP requests via the wireless interface, with one-second pauses in between the completion of a request and the beginning of the subsequent one. The receiver issued 100 requests over TCP without `X-SAP-*` headers to request that the HTTP response be sent on the same TCP socket; then 100 requests over TCP for the same file with an `X-SAP-Force-Precise` header to send the file back via SAP in precise mode; then 100 requests over TCP for the same file to be sent back via SAP in approximate mode. We fixed the bitrate at 54Mbps, the maximum for the 802.11g cards, and measured the total

time from the end of the receiver’s HTTP request to the sender’s precise `SAP_FIN` packet indicating the end of the transmission.

Distance (m)	TCP Time (s)	SAP (Approx) Time (s)	SAP Speedup
2	1.163	0.765	$1.52\times$
4	0.878	0.785	$1.11\times$
6	1.452	0.828	$1.75\times$
8	1.769	1.118	$1.58\times$
10	1.260	0.968	$1.30\times$
<i>Geom. mean</i>			$1.43\times$

Table 1: Speedup of SAP in approximate mode compared to TCP for a 2MB transfer.

Table 1 and Figure 5 show the results of the above experiment with a 2MB randomized binary file over TCP and both modes of SAP. In this experiment, SAP in approximate mode achieved the shortest transfer time at all distances, with an average speedup of $1.43\times$ versus TCP. TCP was the second-fastest transport layer, roughly twice as slow as SAP in approximate mode. At several distances, TCP transfer time increased by nearly two orders of magnitude, skewing the mean transfer time—an unacceptable failure mode for applications that require low inter-packet latency. For this application, SAP’s precise mode used a fixed payload size of 1 KB and did not attempt to implement any form of payload size adjustment, backoff, congestion control, or ACK coalescing, so its slowness relative to the well-tuned Linux TCP stack is not particularly surprising.

Throughput versus bitrate. For this experiment, we fixed the distance between two nodes at 4 m (in the same room), then varied the bitrate across the full range of 802.11b/g bitrates. The sender sent an identical 1 MB file via the wireless interface using each of the three protocols from the previous experiment, ten times per protocol. As before, we sent two initial requests via the wired Ethernet interface to warm the sender’s lighttpd cache. We then repeated the experiment at 6 m and 8 m.

Figure 6 shows the results of the variable-bitrate experiments. When channel quality is good, as it is at 4 m, retransmissions impose little overhead on precise transmissions, so SAP’s savings are not significant. As distance increases, channel quality degrades accordingly, and the precise protocols’ transfer times increase at higher bitrates because damaged frames require retransmission. A worse problem from many applications’ perspective is that the variance of transfer times increases sharply as channel quality drops. SAP in approximate mode exhibits less variance than either precise protocol.

The plots suggest several potential strategies for applications that are faced with falling throughput. Keeping




Figure 5: At a fixed 802.11 bitrate of 54Mbps and increasingly with distance, SAP in approximate mode transfers a 2MB file more quickly than TCP or precise SAP. Black dots represent median transfer times over 100 trials; error bars are centered around mean transfer times and represent standard error.

TCP or precise UDP but switching to a lower 802.11 bitrate is one option, since longer symbols on the air and different coding schemes lessen the impact of interference. However, 802.11 bitrate selection is typically not exposed to applications. An application using SAP can respond in another way, by decreasing the amount of protection applied to payloads. According to these experiments, this alternative approach should allow the application to retain reasonable throughput despite channel degradation.

4.4 Result Quality

To measure the effect of approximate communication on an application with a quality metric, we used the *Tracker* application to compute the cumulative moving average of a walking person’s variable speed, using a recorded trace of 945 data points. The ground-truth average speed, computed directly on the original trace, was 1.56m/s. The cumulative moving average smooths inevitable noise from the GPS receiver and also helps the receiving application smooth incoming readings that arrive with damage.

We set up two SAP nodes as in previous experiments and designated the fixed node as the receiver and the moving node as the transmitter. We fixed the bitrate at 54Mbps and varied the distance between the sender and receiver from 2m to 15m. At each distance, we collected 20 data points with both TCP and SAP in approximate mode, computing the average walking speed.

Figure 7 summarizes the quality degradation we found in this experiment. The median error at each distance was less than 10% versus the ground truth, except at 8m and 15m, where the 18.0% and 83.4% median error respectively represent 0.28m/s and 1.40m/s error in the final speed calculation. At these distances, over half of the values were lost or discarded because of frame errors.

To measure the amount of delay due to retransmissions, we reduced the delay between the sender’s transmissions to 5ms; this makes long delays figure prominently into the total time to transmit the trace. In a track-




Figure 7: For a location-tracking application, varying channel quality (via distance, x-axis) with a fixed bitrate causes output quality (cumulative moving average over a trace, ground truth = 1.56m/s) to degrade gracefully until over half of the values are lost (at 15m). Error bars are centered around the mean and represent standard error over 50 trials per distance.

ing application, long delays may introduce large errors when transmit queues fill with stale data that no longer represents the true location. As in the previous experiments, SAP and TCP achieved similar transfer times under good channel conditions (i.e., at close range). As channel quality degraded, SAP exhibited lower variance, as in other experiments; Figure 8 summarizes.

As described above, the *Tracker* application performs online quality monitoring with a simple sanity check on each incoming GPS reading: if the calculated speed is beyond a reasonable bound, it discards the current distance and speed figures and attempts to interpolate later on. This interpolation tends to overestimate the speed but serves its purpose as a proof of concept. Writing application-specific quality metrics can be subtle [34], but it is a useful exercise beyond approximate computing: any operation that degrades output quality, such as compression, can also benefit.

5 Related Work

SAP is related to other systems that have proposed to allow errors to propagate to applications. In addition to the related work on approximation mentioned in Section 2, several schemes from the networking literature are par-




Figure 6: At a fixed distance of 4 m offering good channel conditions, SAP’s throughput tracks TCP’s. As conditions worsen (at 6 m and 8 m), SAP’s throughput exceeds TCP’s at nearly every bitrate, with less variance. Error bars are centered at the mean over 10 trials for each $\langle \text{bitrate}, \text{distance}, \text{protocol} \rangle$ setting.




Figure 8: SAP exhibits lower variance of inter-packet arrival time than TCP. Error bars are centered around the mean and represent standard error over 20 trials per distance.

ticularly relevant. We focus especially on schemes that question the importance of perfectly precise delivery.

Multimedia streaming. UDP-Lite [24, 23], a variant of UDP with adjustable checksum coverage, *nearly* suffices to provide end-to-end transmission of damaged information. Its designers point out that link-layer checksums—such as the 802.11 FCS—counteract UDP-Lite’s benefits by effectively refusing to carry damaged data; they suggest disabling link-layer checksum checking in device drivers to allow damaged data to pass. SAP builds on this suggestion by coordinating the activities of the link layer and transport layer *and* providing precise communication with complete checksum coverage and guaranteed in-order delivery.

Singh et al. [39] evaluate UDP-Lite for video transmission on cellular links that use the *PPP* protocol at the link layer. They implement a “PPP Lite” mechanism that allows receivers to ignore link-layer checksums and pass damaged data to the application layer. Their technique improved user-perceptible application metrics while streaming video using error-resilient codecs—most notably the inter-arrival time of video frames.

Servetti et al. [38] and Meriaux and Kieffer [27] make similar observations for speech streaming. SAP would be equally useful for the video and audio streaming tasks and also generalizes to other network applications.

Khayam et al. [21] evaluate (in simulation) a SAP-like strategy for multimedia streaming over 802.11b links, with a goal of understanding the distribution of errors that arise as a result of using UDP-Lite. As in SAP, they consider removing integrity checks at low layers to pass potentially damaged data to applications, and they observe throughput improvements with UDP-Lite versus UDP. They advocate using FEC at the application layer, i.e., by encoding content at one end and decoding it at the other, and suggest that UDP-Lite requires less FEC overhead than plain integrity-protected UDP. This technique complements the one we propose for SAP applications; application-layer logic invariably works at a higher semantic level.

Recovering from damage. Some systems have focused on reducing retransmissions by limiting the scope of errors and requesting only small portions of erroneous packets. All of these techniques assume that every bit in every transmission is precious. This assumption underlies the design of even the lowest-level primitives: even though transport-layer protocols like UDP and UDP-Lite [23] can have their checksums disabled or limited, lower-layer 802.11 frames include a mandatory checksum (the *frame check sequence*) that completely cancels any gains from relaxing transport-layer integrity checks.

For some data formats, such as broadcast video, one way to provide best-effort content without retransmission is to transmit the content such that diminished channel quality results in only slight degradations—rather than the choppy playback that conventionally accompanies packet loss. Softcast [19] embodies this approach for video. Redundant coding schemes for other kinds of data (e.g., JPEG2000 images [13] or MJPEG2000 video [14]) to increase error tolerance are an open research area. SAP is independent of such schemes, relying instead on error

tolerance at higher layers; applications may apply their own coding schemes that fit their purposes exactly.

Miu et al. propose *multi-radio diversity* (MRD) [28], which allows a receiver with multiple radios to receive several versions of a damaged frame and uses a selective retransmission scheme instead of 802.11's. MRD's extra radio coordination requires an 802.11 network to be redesigned. SAP is designed for a single-radio scenario and does not have access to information that would allow it to recover damaged frames, but it is also backward compatible with conventional Wi-Fi networks.

Jamieson and Balakrishnan propose *partial packet recovery* (PPR) [20], a method of minimizing retransmissions that allows receivers to use extra confidence information ("hints") provided by the physical layer. With a link-layer retransmission protocol designed to use PPR's extra PHY information, throughput under noisy channel conditions increased by a factor of four. Maranello extends the PPR concept from software radio to commercial 802.11 hardware and can boost UDP throughput by 30% [16]. SAP differs fundamentally in its approach to retransmission; applications need not request retransmission *at all* if the information they receive is "good enough" according to their metrics. To use PPR or Maranello in a manner complementary to SAP, a receiver could lower the confidence threshold it uses to decide whether to request retransmission.

Balan et al. propose TCP HACK [3], a TCP extension that implements TCP header checksums to allow receivers to distinguish between congestion and packet corruption. HACK extends the *partial checksum* idea of UDP-Lite to TCP and requires relaxation of link-layer integrity checks to be useful, so it may be a useful building block toward supporting TCP in SAP. Myriad other proposals attempt to improve TCP throughput under poor channel conditions; we draw attention mainly to HACK because it is the closest in spirit to SAP.

6 Discussion

802.11 and other wireless protocols generally use forward error correction (FEC) in the form of coding schemes with varying efficiency. We are unaware of any WiFi chipsets or drivers that expose such fine-grained control over FEC. If such support existed, SAP could use riskier (more efficient) coding to increase throughput for approximate data.

Non-802.11 networks. A class of devices that may benefit from approximate networking is embedded systems, in which the contrast in power consumption between radio and CPU is stark. For example, a 2.4GHz radio system-on-chip on a sensor mote that is transmitting or actively listening can use nearly $40\times$ more energy than the microcontroller that controls it [12]. On these

systems, reducing the time for which radios are awake is an important design goal.

Recent 802.11 variants. Recent variants of 802.11, particularly the ac variant [29], aggregate many chunks of data per transmission and allow receivers to acknowledge only the chunks they received correctly, an approach reminiscent of partial packet recovery (PPR) [20] and Maranello [16]. Though the SAP prototype is implemented atop 802.11b/g for simplicity, there is no fundamental reason it is incompatible with newer variants. The partial-retransmission behavior of 802.11ac is still predicated on the assumption that the receiver wants to receive packets exactly as they were sent; relaxing this requirement may result in similar gains.

Limitations. A limitation of SAP's approach is that some types of data are fundamentally not error tolerant. For example, SAP cannot support WPA-style encrypted frames, because errors in these frames result in completely incorrect decryption that is unrelated to the correct plaintext. An appealing idea is to redundantly encode ciphertext so that symbols can be recovered despite bit flips, but we assume that decoding these schemes will be prohibitively energy intensive for receivers.

An alternative strategy, in the spirit of QUIC [32], is to support encrypted content is to break it into multiple chunks—e.g., 128-bit AES blocks—and allow for retransmissions only of those chunks that did not decrypt correctly. We leave the design of an appropriate scheme to future work so that it may be explored more fully.

Similarly, some data types are structured such that disturbing any bit will have arbitrarily damaging effects on the output. Highly compressed payloads are an example of such a data type. Other embodiments of approximate computing suffer from similar limitations with respect to encrypted or highly compressed data [35].

7 Conclusion

Approximate computing is a promising technique for better performance and energy efficiency in computer systems. However, to this point it has focused on computation and storage. This paper presented SAP, a cross-layer approach to wireless networking that gives applications a principled way to trade data integrity for better throughput and latency. An application written to use SAP can interleave approximate and precise network transmission modes, and a receiver can elect to accept potentially damaged data and perform its own optional correctness checks.

By showing that optional *approximate* network semantics on a WiFi testbed can improve several important application-layer metrics, we hope that SAP will facilitate exploration of similar relaxations on present and future networks. This effort, combined with prior efforts on

accuracy trade-offs in computation and storage, leads to end-to-end approximate computing opportunities.

References

- [1] AGARWAL, S., MILNER, H., KLEINER, A., TALWALKAR, A., JORDAN, M., MADDEN, S., MOZAFARI, B., AND STOICA, I. Knowing when you're wrong: Building fast and reliable approximate query processing systems. In *SIGMOD* (June 2014).
- [2] AGARWAL, S., MOZAFARI, B., PANDA, A., MILNER, H., MADDEN, S., AND STOICA, I. BlinkDB: queries with bounded errors and bounded response times on very large data. In *EuroSys* (Apr. 2013).
- [3] BALAN, R. K., LEE, B. P., KUMAR, K. R. R., JACOB, L., SEAH, W.-G., AND ANANDA, A. L. TCP HACK: a mechanism to improve performance over lossy links. *Computer Networks* 39, 4 (July 2002).
- [4] BALASUBRAMANIAN, A., MAHAJAN, R., VENKATARAMANI, A., LEVINE, B. N., AND ZAHORJAN, J. Interactive WiFi connectivity for moving vehicles. In *SIGCOMM* (Aug. 2008).
- [5] BORNHOLT, J., MYTKOWICZ, T., AND MCKINLEY, K. S. Uncertain $\langle T \rangle$: A first-order type for uncertain data. In *ASPLOS* (Mar. 2014).
- [6] BRADEN, R. Requirements for internet hosts — communication layers. RFC 1122 (Internet Standard), Oct. 1989.
- [7] CARROLL, A., AND HEISER, G. An analysis of power consumption in a smartphone. In *USENIX Annual Technical Conference* (June 2010).
- [8] DEERING, S., AND HINDEN, R. Internet protocol, version 6 (IPv6) specification. RFC 2460 (Draft Standard), Dec. 1998.
- [9] DÜBEN, P. D., JOVEN, J., LINGAMNENI, A., MCNAMARA, H., DE MICHELI, G., PALEM, K. V., AND PALMER, T. N. On the use of inexact, pruned hardware in atmospheric modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372, 2018 (May 2014).
- [10] ESMAEILZADEH, H., SAMPSON, A., CEZE, L., AND BURGER, D. Architecture Support for Disciplined Approximate Programming. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)* (3 2012).
- [11] ESMAEILZADEH, H., SAMPSON, A., CEZE, L., AND BURGER, D. Neural Acceleration for General-Purpose Approximate Programs. In *International Symposium on Microarchitecture (MICRO)* (12 2012). Selected for IEEE Micro Top Picks 2012.
- [12] FONSECA, R., DUTTA, P., LEVIS, P., AND STOICA, I. Quanto: Tracking energy in networked embedded systems. In *OSDI* (Dec. 2008).
- [13] FRESCURA, F., GIORNI, M., FECI, C., AND CACOPARDI, S. JPEG2000 and MJPEG2000 transmission in 802.11 wireless local area networks. *IEEE Trans. Consumer Electronics* 49, 4 (Nov. 2003).
- [14] GUO, Z., NISHIKAWA, Y., OMAKI, R. Y., ONOYE, T., AND SHIRAKAWA, I. A low-complexity FEC assignment scheme for Motion JPEG2000 over wireless network. *IEEE Trans. Consumer Electronics* 52, 1 (Feb. 2006).
- [15] HAMMER, F., REICHL, P., NORDSTRÖM, T., AND KUBIN, G. Corrupted speech data considered useful. In *ISCA ITRW on Auditory Quality of Systems* (Apr. 2003).
- [16] HAN, B., SCHULMAN, A., GRINGOLI, F., SPRING, N., BHATTACHARJEE, B., NAVA, L., JI, L., LEE, S., AND MILLER, R. R. Maranello: Practical partial packet recovery for 802.11. In *NSDI* (Apr. 2010).
- [17] HOFFMANN, H., SIDIROGLOU, S., CARBIN, M., MISAILOVIC, S., AGARWAL, A., AND RINARD, M. Dynamic knobs for responsive power-aware computing. In *ASPLOS* (Mar. 2011).
- [18] IEEE. *IEEE Standard 802.11-2012*, Feb. 2012.
- [19] JAKUBCZAK, S., AND KATABI, D. SoftCast: Clean-slate scalable wireless video. In *Allerton Conference on Communication, Control, and Computing* (Sept. 2010).
- [20] JAMIESON, K., AND BALAKRISHNAN, H. PPR: Partial packet recovery for wireless networks. In *SIGCOMM* (Aug. 2007).
- [21] KHAYAM, S. A., KARANDE, S., RADHA, H., AND LOGUINOV, D. Performance analysis and modeling of errors and losses over 802.11b LANs for high-bit-rate real-time multimedia. *Signal Processing: Image Communication* 18, 7 (Aug. 2003).
- [22] KNESCHKE, J. Lighttpd. <http://www.lighttpd.net/>.
- [23] LARZON, L.-A., DEGERMARK, M., AND PINK, S. UDP Lite for real time multimedia applications. In *IEEE International Conference on Communications (ICC)* (June 1999).
- [24] LARZON, L.-A., DEGERMARK, M., PINK, S., JONSSON, L.-E., AND FAIRHURST, G. The Lightweight User Datagram Protocol (UDP-Lite). RFC 3828 (Proposed Standard), July 2004.
- [25] LIU, S., PATTABIRAMAN, K., MOSCIBRODA, T., AND ZORN, B. G. Flickr: Saving refresh-power in mobile devices through critical data partitioning. In *ASPLOS* (Mar. 2011).
- [26] MARIN, C., LEPROVOST, Y., KIEFFER, M., AND DUHAMEL, P. Robust MAC-lite and soft header recovery for packetized multimedia transmission. *IEEE Trans. Communications* 58, 3 (Mar. 2010).
- [27] MERIAUX, F., AND KIEFFER, M. Robust IP and UDP-Lite header recovery for packetized multimedia transmission. In *IEEE Intl. Conf. Acoustics Speech and Signal Processing (ICASSP)* (Mar. 2010).
- [28] MIU, A., BALAKRISHNAN, H., AND KOKSAL, C. E. Improving loss resilience with multi-radio diversity in wireless networks. In *Mobicom* (Aug. 2005).
- [29] ONG, E. H., KNECKT, J., ALANEN, O., CHANG, Z., HUOVINEN, T., AND NIHTILÄ, T. IEEE 802.11ac: Enhancements for very high throughput WLANs. In *IEEE Personal Indoor and Mobile Radio Communications (PIMRC)* (Apr. 2011).
- [30] POSTEL, J. Internet protocol. RFC 791 (Internet Standard), Sept. 1981.
- [31] RIEMANN, R., AND WINSTEIN, K. Improving 802.11 range with forward error correction. Tech. Rep. MIT-CSAIL-TR-2005-011, Massachusetts Institute of Technology, Feb. 2005.
- [32] ROSKIND, J. Experimenting with QUIC. <http://blog.chromium.org/2013/06/experimenting-with-quic.html>, June 2013. Visited September 21, 2014.
- [33] SALTZER, J. H., REED, D. P., AND CLARK, D. D. End-to-end arguments in system design. *ACM Trans. Comput. Syst.* 2, 4 (Nov. 1984).
- [34] SAMPSON, A., DIETL, W., FORTUNA, E., GNANAPRAGASAM, D., CEZE, L., AND GROSSMAN, D. EnerJ: Approximate Data Types for Safe and General Low-Power Computation. In *Conference on Programming Language Design and Implementation (PLDI)* (6 2011).
- [35] SAMPSON, A., NELSON, J., STRAUSS, K., AND CEZE, L. Approximate storage in solid-state memories. In *International Symposium on Microarchitecture (MICRO)* (12 2013).
- [36] SANDVINE INC. Global internet phenomena report. <https://www.sandvine.com/downloads/general/global-internet-phenomena/2014/1h-2014-global-internet-phenomena-report.pdf>, May 2014. Visited September 29, 2014.

- [37] SEN, S., GILANI, S., SRINATH, S., SCHMITT, S., AND BANERJEE, S. Design and implementation of an “approximate” communication system for wireless media applications. *SIGCOMM Computer Communication Review* 41, 4 (Aug. 2011).
- [38] SERVETTI, A., AND DE MARTIN, J. C. Error tolerant MAC extension for speech communications over 802.11 WLANs. In *IEEE Vehicular Technology Conference (VTC)* (May 2005).
- [39] SINGH, A., KONRAD, A., AND JOSEPH, A. Performance evaluation of UDP-Lite for cellular video. In *Intl. Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)* (June 2001).
- [40] STONE, J., GREENWALD, M., PARTRIDGE, C., AND HUGHES, J. Performance of checksums and CRC’s over real data. *IEEE/ACM Trans. Networking* 6, 5 (Oct. 1998).